Gmail **John Loop <pccitizen@gmail.com>**

John Loop's 2-1-2023 imonitor[g] newsletter. DNS/Speedtests. This project may never end.... Work continues on refining scripts for 2nd,3rd,4th order effects!

John Loop <jdloop@imonitorg.com>
To: "pccitizen@gmail.com" <pccitizen@gmail.com>

Dear imonitor[g] users, potential users, former users, and interested parties:

*Summary for those of you new to this newsletter:**
I am working on a project to perform Internet/Local Network monitoring. I use a small device[raspberry pi3B or pi3B+], along with custom scripts partners across the country representing many different access technologies and ISPswho help in the development. I initially targeted the servic universally, and I have helpers across the country on Windstream, ATT, Ting, Spectrum, Comcast, CenturyLink, Tmobile, Starlink and several oth VDSL, WADSL, cellular,cable, satellite and fiber. I also released a "generic" version which iscompletely standalone, and has no connection to m on sourceforge.net. I also established a github page pointing to it. These images run standalone and contain no links to my server, nor do I hav third party server in the internet which you have to access to collect your info --it is all on a webpage on your local raspberry pi! All you have to d it from me if you want], plug it into your raspberry pi3B or 3B+, connect the ethernet and enjoy. You can configure an email using your gmail acc https://imonitorg.com

This is the perfect use for an older raspberry pi 3B or 3B+ many of you have lying around! And I hope to transition the scripts to a pi4 this year.

For my 25 trial users, I continue to appreciate the use of your ISP. Thomas from Stockton CA has joined -pi16, so I now have an ISP [comcast] opposed to a generic/anonymous user], I have pi5 and pi30 waiting to be reassigned [I provide the raspberry pi!]. All you have to do is give me when you receive your rpi. You will receive a daily email detailing the performance stats and plots of your ISP connection [I will preconfigure you listed. Most of the features are configurable on the local web page. Trial user's rpi's are accessed by me for testing and update, unless you spe pis operate like this [pi16,29,31]. Rpi16,29 and 31 [and all the other rpis] report their stats to a web page I maintain for all these rpis:

https://imonitorg.com/customerplots/customerplots.html

The main page contains lots of additional information: https://imonitorg.com

2-1-2022 Newsletter - Several topics have consumed me since the last newsletter.

1. Current partners, as a recap. A good range, but would like some more!! Of course this doesn't count the "generic" pis which are not "manage

-comcast cable: pi3, pi25, pi26 (Denver metro, Joey,Ken, Fred), pi16 (Stockton CA, Tom), pi14 (Atlanta metro, Matt)

-TDS Telecom: pi32 (Denver metro, Fred)

-ATT fiber: pi22 (Atlanta, Charley), pi23 (Dallas metro, Kenny)

-Spectrum cable: pi5,10,11,15 (Tampa FL metro John, Mike), pi21 (Dallas metro, Jerry)

-ATT Uverse: pi7,13,12 (Bishop GA, Phil) pi28 (Ptree City GA, Kevin)

-Starlink: pi8 (S GA. Mel), pi9 (N GA, Ernie)

-Ting Fiber: pi24 (Raleigh NC metro Niel)

-Longmont Fiber: pi27 (N CO, Lou)

-PentelData: pi6 (N New Jersey, Rick)

-WADSL (Vistabeam, FrontRange): pi29, pi31 (N CO Tom)

-Spectrum mobile home broadband [pi1,20] (Tampa FL metro John)

Recently retired:

VDSL: pi3,9 (N GA)

ATT Uverse: pi30 (Chicago metro Bob)

Hughes: (N GA)

Tmobile Home broadband [Raleigh metro cellular broadband]

2. DNS oddities/complexities

DNS continues to be the -almost mysterious- underlayment to everything that we do on the Internet. It is very difficult to detect exactly what is ha "Internet" could exist without DNS of course [that is starting to be somewhat of an exaggeration with the "tricks" employed these days]! So it is a

are no "connections" to track. And it is largely in clear text, though there are transitions to encrypted techniques.

In a normal home LAN, your PC/etc receives a DNS server IP address from your router via DHCP. This DNS server is used for all initial DNS qu within the browser, within the router, and typically in the -ISP- DNS server, all the rest, in line until the "authoritative" DNS server is contacted. An to live, which can vary. Even "failures" are cached as "NXDOMAIN" -non existent domain. you may see occasionally when you mistype a doma experienced times when you have a bad cache and you need to clear it. There are ways to do this in linux/windows/browsers -but this is not con

Complicating things further, browsers are now VERY stubborn about allowing anything but https! You have to navigate and ALLOW unencrypted

To complicate the DNS matters, all major browsers now do DNS queries [and caching] from "within" the browser using a browser assigned DNS "hides [encrypts]" the DNS query [at least the first query [the remaining recursions?], and the DHCP assigned DNS server is not used. You can c provider. Currently, only NON-browser activity on your PC/etc uses the DHCP-assigned DNS server, which is still a "lot" of stuff. Just think of all on. You have almost no conception of the extent of this activity.

If your house is flooded with "IOT" gadgets these days [like mine], you should make sure you put them all on a "guest" network on your router. T connecting to servers all over the world checking for updates/syncs/ads/etc, and passing info about your IOT gadget and your DNS activity. The well as the questionability of the integrity of the IOT code. Much of this code is old linux-based and never updated. Older cameras and smart TV network! This is the source of most of the botnets in the world -compromised IOT gadgets. The DNS activity of these gadgets is equally vast!

The DHCP-assigned DNS server is usually the router itself, which does DNS proxying/forwarding and caching. You can override this by specifyi on your router, which would enable you to change the DHCP-provided DNS server value. The manipulation/nulling of DNS queries is the technic and is provided on the rpi. To use pihole or an equivalent potential ad-blocker on your LAN you need to steer DNS queries to the pihole server. queries to the rpi and actually see the DNS activity. You will be amazed. If you want, you can even subscribe to lists which will do massive ad bl reports DNS queries.

The DNS server IP that the router uses on the WAN side is sometimes configurable, but is usually the ISP-provided DNS server. ISPs often inte encrypted]. It is a record of your DNS queries and your IP address -very useful info to ad companies. In this day and age, the IP address assigr very useful. Your IP address is almost like a house address. However, with the advent of DOH, this data is not interceptable [and the DOH prov decrypted on their endpoint]. One way to minimize this compromise is to specify a generic DNS server where you can. There are four main one 208.267.220.123 opendns. They are located "near" because of anycast and CDN networks, so there is little compromise wrt to speed in using th your ISP DNS server!! Some of them filter for malicious sites and offer ways to filter.

So the DNS activity is a chain of UDP requests [recursion] from your PC into the Internet [and the resultant DNS responses], very difficult to tracl typically located a few hops into your Internet, and they synchronize "on the back end" to "authoritative" DNS servers. In the old days, the updat hours, but these days, the DNS update is very quick, especially with the major domain providers.

You can bypass this UDP chain by providing a specific DNS server IP in your PC, such as using "1.1.1.1 [or 8.8.8.8 or 9.9.9.9 etc]" in your DNS directly to 1.1.1.1, bypassing the caching/relaying/proxying using normal provided DNS servers. The downside is that these DNS servers -may- use. They are still doing recursion if necessary, but they have such massive caches that the recursion is rare.

The DNS server in your LAN will perform "recursive" lookups, and queries will be repetitively issued to the domains on the Internet until it reache Your router normally performs this [if the value is not cached], so every lookup is visible. There are ways to designate servers further upstream t

There is a popular site which "benchmarks" DNS servers, but this is almost overkill today https://www.grc.com/dns/benchmark.html

I would very much welcome any comments/criticisms/corrections wrt to my DNS understands!!! Please!

2. DNS related testing in the rpi scripts

The rpi uses the DNS server assigned via DHCP in the user's network to do DNS queries when using tcp pings. This is by design since we mean

The tcp pings use names -a "top100web" sites list. The tcp ping is done via DNS FQDN [fully qualified domain name] to invoke a DNS query us some measure of performance of the user's network.

Of course the primary ping is the ICMP ping -using IP address only- to a "near target" -thus gauging Internet performance regardless of the DNS the daily graph.

Two DNS related "pings" are performed twice per hour -plotted twice per hour. One is a DNS query of www.google.com directed to the home rou home network and the home router [I first verify that the home router is listening on the DNS port]. Most of the time, the DNS query to the home response time is near 0. This is the plot of the solid circles twice per hour. You can see when the router has to occasionally relay to an upstrean

The second "ping" is a curl of www.google.com [configurable in this case] which "represents" the highest layer of "ping" -simulating a browse to a plot, and typically have the maximum delay of the plotted variables. You can often see these "bubbles" rise during the aft/evening as congestion

3. DNS techniques in the scripting

I go to considerable depth to not fall prey to network failures or inconsistencies. I use actual IP addresses such as 1.1.1.1 in the early tests to ch any home network DNS-related variable before I rely on LAN DHCP values. This is all before I start the routine pinging. I have these actual IP a Nevertheless, understanding and quantifyng DNS performance is very difficult, and I continue to experiment.

4. A new network analysis package.

I have recently discovered the open source project "ntopng," which is a "next generation" version of the popular linux utility "top." I have installed ***NOT installed the package*** [all of you who are trial participants]. With the addition of ntopng, the user has the ability to perform a higher layer a knowledge above that acquired by wireshark. This package is under active development [free and paid versions of course]. There is still the ne acquired by ntopng. This will no doubt be coming in the years to come. There is a desperate need for a package to digest all the activity on a us

Right now, all we can get is INFORMATION about what is going on and interpret it ourselves. An AI can actually infer important information abou interpretation. This will be crucial in coming years to ID malware/malicious activity. Ntopng is a step on that road.

Unfortunately, ntopng is so resource intensive that its support has moved to 64 bit architectures only, which leaves raspberry pi in the dust with tl a stretch pi3B+, but it seems to break. Installing it on a buster pi4B works much better, but this is still only 32 bit. Stay tuned for more developm be very valuable, which is why I have a "hook" embedded for future tests.

In a more general sense, you can install ntopng on windows or MAC, or on some distributions of Linux. You should try it. This is a wonderful too pihole you are flooded with information. I have installed it on a separate rocky linux PC. As I said, this is good and bad. So much info, so little a with useful network tools. Go to https://www.ntop.org/products/traffic-analysis/ntop/ Here is a recent article as well: https://www.admin-magazin ntopng

5. Pihole

Pihole has been installed on all the rpis for about a year now. As an interim technique to reveal what is going on with your network, this goes a l page to this software. One to the pihole dns log, which records all DNS activity encountered by the rpi. By pointing your PC/etc DNS server to tl off DOH in your browser to see "everything"]. There is also a link to the actual pihole app, where you can install ad filters if want. You will be abs your PC/etc. You will also be amazed at all the ipv6 activity in the background. This screenshot of the rpi shows the quick links to important data "near real time" plot of performance [discussed last newsletter -below]."

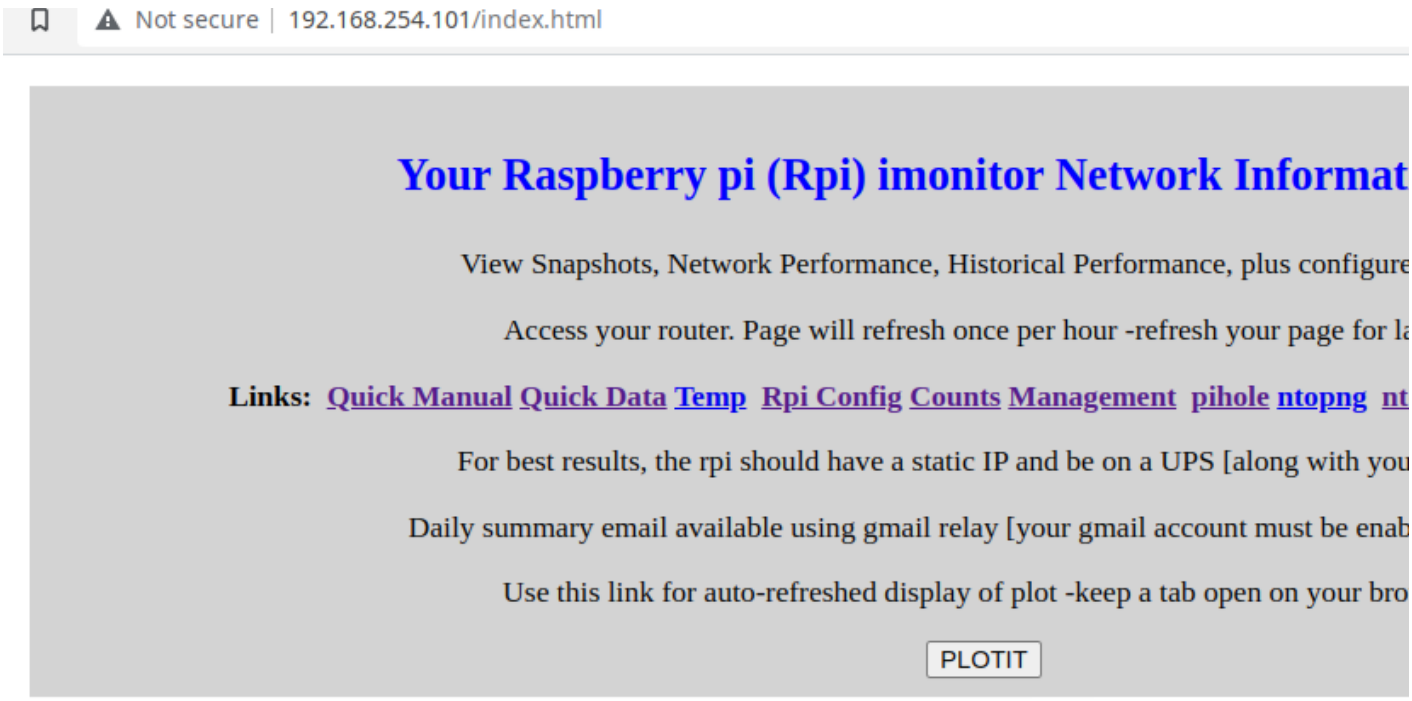

6. Daily email for participants

The daily email is getting pretty clogged with information, and I am looking at ways to condense the information. I need an "AI" to interpret all thi of hosts on their networks, the host reporting can be especially intimidating, listing all the hosts that come and go, and are "persistent." You shou listed on the web page, including an historical arp table, which shows all the devices which have EVER beeb on your network [with a "clear" opti

7. SDcard reliability

The longterm viability of the microSD card continues to be an issue. These microSD cards were not meant for OS intensive tasks, in general, bu "writing" which wears out the microSD card. I have developed a script [plus alert email] which I hope will detect when a microSD cards enters fa this is not a sure thing. I can often "recover" the SDcard once it fails by simply reading it in my linux machine and writing to a new SDcard! I hav

8. Realtime performance plotting

I have improved the "near real time" performance plotting by giving you a way to make it update every 1 minute, instead of the default 5 min. Th plot. If you have not used this, I would highly recommend it. You get a link to it in each daily email. It will occupy a separate tab in your browser

9. Speedtests and interpretation

Over the past couple years, nothing is more "confusing" than speedtest results. Determining the "actual" speed on a statistically multiplexed net the providers use? I seriously doubt it. If I get two back to back packets can I seriously say that the speed is the same as the "line" speed? "Lin customer has an individual "line" to their router. Aggregation occurs in the ISP. Cable/satellite is a different matter! EVERYBODY is on that "line has their "own" speedtest app or website, and it is of course optimized for their network [surprise]. The pi uses the linux code for the "speedtest. server. I don't know how this algorithm works. What I do know is that some speedtest servers are much better than others, let alone that they m provider's network. I also know that speedtests vary considerably from moment to moment and day to day. And this will CONTINUE to be the c

speed, networks that run at infinite speed, and no buffering is used anywhere. That will never happen of course. The ONLY variable in this setu photons in their medium]. TCP algorithms adjust to the delay with very complex algorithms, variously implemented on different platforms!

The pi has a historical graph showing the speedtest results, and this is useful for indicating long term trends. You may be encouraged to power change -somebodyis DOSing your IP address. You may need a new router [all these situations I have seen!]. You may also see gradual degrad factors.

There are many mysteries in determing/collecting/interpreting speedtest results and I am far from figuring them out. I HAVE discovered a particu much more closely to "line" speed than others. Using normal servers selected by the speedtest algorithm have been UNABLE to reveal the upst and a few others. It is very frustrating to see the speed for a 300Mbs/300Mbs fiber being reported as 300/30 consistently. Specifiying the server competent. This is a speedtest mystery I have yet to figure out [and may never]. It is as if the TCP algorithm is optimized when they "talk" to ea

In the interim, I find that the speedtest server "speed.cloudflare.com" very revealing. It plots your response which shows how the speed varies d MANY variables in the TCP algorithm which contribute to the performance, as well as the intervening layers of routers and switches. I believe th protocol [see previous newsletter] which may make a difference.

9. Examples

I have included inline the last newsletter because it shows a lot of great examples of plots available.

10. Continuing

I continue to appreciate the participation of my partners in this adventure. I am especially open to suggestions and criticisms for improvements.

John

The last newsletter 7-2022:

Consider this an interim newsletter. Wonderful example plot showing problems. Starlink perf plots as Ernie and Mel adjust their sats. Oth

Example 1 Ernie in N GA switches from Windstream to Starlink

The first attached perf plot for Ernie shows when he removed his pi3 from Windstream and rebooted it on Starlink about 4PM [1600]. You then note the change in the baseline response. There is an additional 50-100 msec to get up/down from the Starlink sats. Plus, Ernie stil of timeouts in the plot. He still has trees "up top." The "packet loss" reported in the speedtests is maybe 4%, high, but the connection is much better than with Windstream. Ernie's windstream connection is probably the worst of my trial people, mostly because he has a milli this date, the up/down on Windstream VDSL [20Mbs/2Mbs] is pretty slow compared to most participants these days.

*Example 2 Mel in S GA Starlink looks good *

The second attached perf plot for Mel on pi8 shows a Starlink connection after he has roof mounted it. It is "almost normal" in that there i at the "elevated" level. His early perf plots looked like Ernie's starlink connection above, because he had the sat antenna on the ground.

Example 3 Niel Perfect on fiber in NC

The third attached plot is a "perfect" perf plot. It shows Niel's pi24 Ting fiber connection in Rolesville NC. Niel started out with CenturyLin gadgets trying to "check-in/sync'], went to Tmobile [cellular 50Mbs], and is now on Ting [about 300/300Mbs up/down]. He went from the year. Lucky Niel!

Example 4 A speedtest archive plot from pi4 showing change of router

Joey recently purchased a new router for his pi4 comcast connection in Englewood CO. You can see the difference the new router made typical of cable. It shows a wide variation consistent with the shared medium of the cable. This speedtest archie is available for all pis on

*Example 5 Repeated from 4-10 newsletter because it is so instructive. Explains perf plot. *

My friend Tom in North CO uses Front Range Internet ISP on pi29, a wireless ISP in North CO. "Wireless" means "wireless ADSL" not ce and has to worry about things like birds camped out on his antenna! They have been having terrible problems with their service the past some very interesting information. I think it is a great example of how useful a monitor service like this can be to gauge Internet/local netw

As a summary before discussing the plot, the performance plots generated by the imonitor scripts running on the pi3B[+] all show 8 things

---beware, the colors may differ from what the attached plot shows.

1. "near" ICMP ping response times [IP address] -results of once/min ping -purple? first line typically 5 hops or less, tho Tom has his naile -these are the spikes to 1000 ms [1 sec].

--in addition, if there is a near ICMP ping timeout, 6 ADDITIONAL [hyper]pings are attempted WITHIN that minute to gauge "time-depth" ticks [crosses] which appear on the spikes, starting at 500ms and going to 620ms if all 6 timeout.

2. tcp half open SYN responses [SYN, wait for SYN/ACK and then send FIN] using top 100 web sites [this usually goes to the nearest CD -green? second line

3. "far" ICMP ping response times [near ICMP targets of other trial members -about 20 addresses spread around country between diff ISP

4. tcp timeouts [when 2 above shows a timeout] - yellow? solid boxes at 220ms

5. far ICMP timeouts [when 3 above shows a timeout] -blk? triangles at 320ms

6. DNS query using local router -red solid dots [using local router as DNS cache/proxy] - no response within "1000 ms" is declared a DNS

7. Curl of Internet website [www.google.com in this case] -clear dots [simulated web page draw] -no response within "1000 ms" is declare

8. HARD "Offlines" -which uses an algorithm of near ICMP timeout+tcp timeout -red? clear boxes at 450ms

The HARD "offline" is a judgement call, using my own algorithm. In reality of course, any INSTANT of time could be an Internet offline. I a near ICMP failure which crosses a tcp failure in a one minute interval, and each msmt having a timeout declared as "no response in 1 mir 10AM and around 12PM. Completely arbitrary definition of "offline," but which hopefully may represent what the user "sees."

This is a daily plot, generated at 1AM Sunday morning for the previous day [actually 2AM yesterday until 1AM today -Sunday]. Tom receiv additional network information. This information is also accessible via a web page running on the raspberry pi at any time.

Looking at Tom's pi29 plot now, which is suggestive of a lot of problems, you can see that it is "stable from 2AM to 9AM, has some proble down" about 2PM [1400], as if they have the right network configuration/routes, etc. Maybe the network engineers thot things were fixed b gradually get worse starting about 3PM [1500] until it looks BAD about 7PM [1900], and at 8PM [2000] things are looking so crummy that configurations again. The plot from here looks completely different, as the near ICMP pings are failing now [and the 6/min hyperpings -pr target], but the tcp pings are "almost" returning to normal, as well as the DNS queries and the curls. "Near normal" service looks possibl have disappeared/been replaced? in the network reconfiguration [showing the near ICMP and hyperping failures], but the tcp looks a lot b

In spite of all these problems, there is no persisting HARD offline declared during this entire interval, and not after the "reset" at 8PM until apart. So "service" from 8PM to 11PM is "normal?" -looking beyond the near ICMP failures, which may be using targets gone/not in the p

At about 11PM everything goes to hell again tho, which continues to about 1230AM [wrapping to left of plot] [Plot BEGINs at 2AM on left, on left 1AM to 2AM is NOT plotted]

Any further interpretations would really be welcome. I spend a lot of time trying to manipulate these scripts, plots and algorithms to revea

Thanks to all my trial people for letting me do this! You could add one of these monitors to your network with just a little work!

My next newsletter will detail the daily email and all that it shows.

John

--
https://imonitorg.com

5 attachments

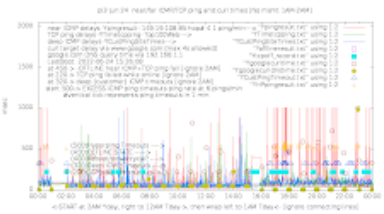

pi3Jun25specTOstarlink.png
33K

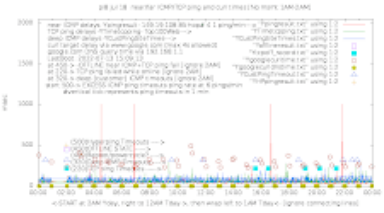

pi8Jul19AlmostNormal.png
30K

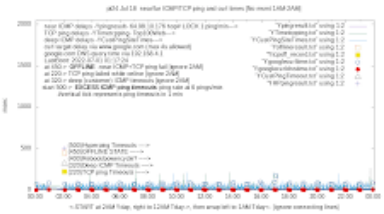

pi24Jul19best.png
28K

pi4-SpeedtestArchiveJul18.png
19K

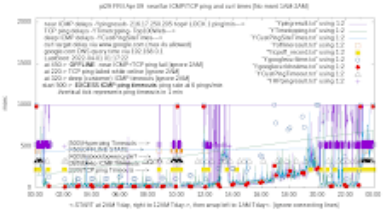

GREATdebuggingplot4-9-2022.png
35K